

*RESPONSE RATE VIEWED AS ENGAGEMENT BOUTS:
RESISTANCE TO EXTINCTION*

RICHARD L. SHULL, SCOTT T. GAYNOR, AND JULIE A. GRIMES

UNIVERSITY OF NORTH CAROLINA AT GREENSBORO

Rats obtained food pellets by nose poking a lighted key, the illumination of which alternated every 50 s during a session between blinking and steady, signaling either a relatively rich (60 per hour) or relatively lean (15 per hour) rate of reinforcement. During one training condition, all the reinforcers in the presence of the rich-reinforcement signal were response dependent (i.e., a variable-interval schedule); during another condition only 25% were response dependent (i.e., a variable-time schedule operated concurrently with a variable-interval schedule). An extinction session followed each training block. For both kinds of training schedule, and consistent with prior results, response rate was more resistant to extinction in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. Analysis of interresponse-time distributions from baseline showed that differential resistance to extinction was not related to baseline differences in the rate of initiating response bouts or in the length of bouts. Also, bout-initiation rate (like response rate) was most resistant to extinction in the presence of the rich-reinforcement signal. These results support the proposal of behavioral momentum theory (e.g., Nevin & Grace, 2000) that resistance to extinction in the presence of a discriminative stimulus is determined more by the stimulus–reinforcer (Pavlovian) than by the stimulus–response–reinforcer (operant) contingency.

Key words: resistance to change, behavioral momentum, extinction, bouts, visits, key poke, rats

Much of the evidence that supports behavioral momentum theory (Nevin, 1992; Nevin & Grace, 2000) has come from the following kind of procedure. Two discriminative stimuli are presented in an alternating sequence during training sessions. One of the stimuli signals a relatively rich variable-interval (VI) schedule (e.g., VI 1 min), and the other signals a relatively lean VI schedule (e.g., VI 4 min). After many such training sessions, conditions are altered so as to reduce response rate and thereby assess the relative resistance of the response to change in the presence of the two discriminative stimuli. For example, following training, the procedure might be changed to extinction in the presence of the two discriminative stimuli. The highly reliable finding is that response rate declines more slowly, relative to its baseline level, in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal (see Nevin, 1992; Nevin & Grace, 2000, for reviews).

This research was supported by grants from the National Science Foundation (IBN-9511934 and IBN-0125093). Scott Gaynor is now at Western Michigan University.

Correspondence concerning this paper should be directed to R. L. Shull, Department of Psychology, Box 26164, University of North Carolina–Greensboro, Greensboro, North Carolina 27402-6164 (e-mail: rlshull@uncg.edu).

It appears, moreover, that resistance to the decremental effects of extinction and other variables depends not so much on the relative reinforcement of the designated response in the presence of the discriminative stimulus but rather on the relative reinforcement regardless of responding in the presence of the stimulus. That is, resistance to change seems to depend more on the stimulus–reinforcer (i.e., Pavlovian) contingency than on the stimulus–response–reinforcer (i.e., operant) contingency. Evidence favoring this conclusion comes from procedures in which, during training, additional reinforcers are presented independently of the designated response in the presence of one of the discriminative stimuli (Dube, Mazitelli, Lombard, & McIlvane, 2000; Grimes & Shull, 2001; Mace et al., 1990; Mauro & Mace, 1996; Nevin, Tota, Torquato, & Shull, 1990). For example, one of the discriminative stimuli might signal a VI 4-min schedule and the other might signal a VI 4-min schedule plus a concurrently available variable-time (VT) 1.33-min schedule. The total rate of reinforcement during the latter discriminative stimulus, then, is about four times higher than that during the former discriminative stimulus. But the rate of response-dependent reinforcement (for the designated response) is the same for the two discriminative stimuli. The result is that re-

sponse rate during training (i.e., baseline) is either lower in the presence of the stimulus that signals the VI + VT schedule than in the presence of the stimulus that signals the VI schedule alone or is about the same in the presence of the two discriminative stimuli (Burgess & Wearden, 1986; Rachlin & Baum, 1972). Yet responding is reliably more resistant to change—to extinction, for example—in the presence of the stimulus that had signaled the overall higher rate of reinforcement (i.e., the VI + VT schedule). Thus, neither the relative reinforcement of a designated response nor the relative rate of that response during training predicts the relative resistance of the response to change. Instead, relative resistance to change seems to depend largely on which discriminative stimulus has signaled the higher rate of reinforcement regardless of responding. That is, it depends on the stimulus–reinforcer (Pavlovian) contingency.

Although the results cited in support of this conclusion are reliable, questions have been raised about their generality. A potential limitation is that the performance measure for assessing resistance to change (i.e., response rate) is typically computed by treating all instances of the designated response as equivalent. Yet reinforced responding probably is better viewed as organized into bouts of responses (or visits) separated by periods of disengagement (Blough, 1963; Davison & Charman, 1986; Gilbert, 1958; Mellgren & Elsmore, 1991; Nevin & Baum, 1980; Pear & Rector, 1979; Premack, 1965; Shull, Gaynor, & Grimes, 2001). Moreover, the rate of initiating bouts and the responding within bouts are controlled by different classes of independent variables (Blough, 1963; Mechner, 1992; Shull *et al.*, 2001). Thus, response rate, as typically computed, is a composite measure. Variables that affect total response rate similarly can affect the components of response rate (i.e., bout-initiation rate, response rate during bouts, and the length of bouts) differently. Bout-initiation rate, for example, is consistently related to relative reinforcement variables even when total response rate is not (Shull *et al.*).

Conceivably, then, the interpretation of a resistance-to-change test could depend on whether the performance measure is total response rate or bout-initiation rate. For ex-

ample, the different rates of baseline responding under a multiple VI 1-min VI 4-min schedule seem to result mainly from corresponding differences in bout-initiation rate (Shull *et al.*, 2001). Assuming that such findings for bout-initiation rate under VI schedules are general, differential resistance to extinction following training under a multiple VI 1-min VI 4-min schedule would be correlated with the differential disposition to initiate bouts of responding established during baseline. If conditions were changed such that 75% of the 60 reinforcers per hour during the rich-reinforcement signal were now delivered independently of responding (i.e., a VI 4-min + VT 1.33-min schedule), the response rate would be lower than when all the reinforcers were response dependent (i.e., during the former VI 1-min component). This lower total response rate in the rich-reinforcement component might result from a relatively high bout-initiation rate combined with a relatively low number of responses per bout compared to the rich VI component alone. Thus, unlike the multiple VI 1-min VI 4-min schedule, a multiple VI 4-min + VT 1.33-min VI 4-min schedule might engender similar total response rates across components during baseline training. But the disposition to engage in the operant behavior, as measured by bout-initiation rate, might remain higher in the presence of the VI 4-min + VT 1.33-min signal (as seems likely for the rich VI component of a multiple VI VI schedule). If that happened, resistance to change in both the multiple VI (rich) VI (lean) and multiple VI + VT (rich) VI (lean) schedule situations would, in fact, correspond to the baseline rate of operant behavior, in the sense of bout initiations rather than individual instances of the designated response. However, the similarities in total response rate during the baseline training with VI + VT (rich) and VI (lean) would have masked the systematic differences in bout-initiation rate, which may be critical for producing the differential resistance to change. Following this line of reasoning, we compared total response rate and bout-initiation rate using both multiple VI (rich) VI (lean) and multiple VI + VT (rich) VI (lean) schedules.

There are, as well, other possible combinations of different bout-initiation rates and responses per bout that could result in similar

total response rates between the two discriminative stimuli during baseline. In addition to different bout-initiation rates, the two kinds of schedule (VI and VI + VT) might generate different bout lengths (i.e., different average numbers of responses per bout), and these response units might be differentially susceptible to disruption by extinction. Differential resistance to extinction at the level of total response rate might better be interpreted, then, as reflecting the stability of different operant units (e.g., different bout lengths) formed during training than as reflecting the relative strength of equivalent operant units. Other researchers appear to have had similar concerns (Branch, 2000; Galbicka & Kessel, 2000; Shimp, 2000; see also Doughty & Lattal, 2001, and Neuringer, Kornell, & Olufs, 2001, for additional evidence that different operant units can be differentially susceptible to disruption).

A different approach to the same issue emphasizes possible parallels between the pause–bout character of reinforced responding and features of natural foraging behavior. Mellgren and Elsmore (1991) suggested that the pauses between bouts might be analogous to the search mode in foraging, and responding within bouts might be analogous to procurement or prey handling. From this conceptualization, and from considering data from studies of foraging, Mellgren and Elsmore were led to the interesting hypothesis that the relation between resistance to extinction and the rate of reinforcement during training might be different—in fact opposite—for responding within bouts and for bout initiations. They reasoned that the relation would be positive for within-bout responding (consistent with the results described above on resistance to change) but negative for bout initiations. Indeed, they found some evidence for this sort of dissociation in the performance of rats that obtained food by lever pressing and then were subjected to extinction. For this analysis, in their study, rate of reinforcement was varied between different groups of rats; each rat was trained under a single rate of reinforcement. It is therefore unknown whether similar results would be obtained following training in which two different discriminative stimuli alternate during a session, each correlated with a different rate of reinforcement. Such pro-

cedures are likely to be critical for establishing the sorts of stimulus–reinforcer contingencies that appear to determine resistance to change, at least as assessed by total response rate (Cohen, 1998). Yet from Mellgren and Elsmore's results and hypothesis, one might expect bout-initiation rate to be *most* resistant to extinction in the presence of the discriminative stimulus that signaled the *lower* rate of reinforcement during training, contrary to the basic generalization of behavioral momentum theory (Nevin & Grace, 2000).

There are, then, grounds for supposing that resistance to extinction might be positively correlated with bout-initiation rate during baseline, even when only some of the rich-schedule reinforcers are response dependent. At the same time, there are grounds for supposing that the correlation might be negative (Mellgren & Elsmore, 1991). Either outcome would suggest limits on the conclusion that resistance to extinction depends more on the stimulus–reinforcer contingency than on the response–reinforcer contingency.

The purpose of the present study was to assess the possible relations between bout-initiation rate and bout length on the one hand, and resistance to extinction on the other. This assessment was carried out in three steps. The first involved replicating, with total response rate as the dependent measure, the previously reported finding that responding is most resistant to extinction in the presence of whichever discriminative stimulus has signaled the higher rate of reinforcers, whether those reinforcers were delivered by a VI schedule or by a VI + VT schedule. The next step involved examining the relations between baseline measures of bout performance (i.e., bout-initiation rate and bout length) and resistance to extinction. Finally, the resistance to extinction of bout-initiation rate and bout length was assessed.

Rats obtained food pellets by nose poking a key that was illuminated either with a blinking light or with a steady light. These light-status components (i.e., the discriminative stimuli) alternated throughout each session and signaled either a relatively rich or a relatively lean rate of reinforcement. For one training condition (i.e., a block of training sessions), all reinforcers were response dependent. That is, the discriminative stimuli

were correlated with different VI schedules, one rich and one lean. For another training condition, approximately 75% of the reinforcers in the presence of the rich-reinforcement signal were delivered independently of responding (i.e., VI + VT). An extinction test was given following each of these training conditions. The novel contribution of our study comes from recording interresponse times (IRTs) during baseline and extinction so as to assess the effects of the training schedules and extinction on bout-initiation rate and on the average number of responses per bout.

METHOD

Subjects

Eleven male Long Evans hooded rats, identified as members of Squad 1 (4 rats), Squad A (4 rats), or Squad B (3 rats; a 4th rat in this squad died before the start of the present project), served as subjects. The rats were housed in individual plastic cages offering free access to water. The temperature in the housing room was maintained at about 22 °C, and the overhead lights in the room were on from about 6:00 a.m. until 6:00 p.m. each day; the experimental sessions were conducted during the lights-on periods.

The rats were obtained from a commercial breeder at about 6 weeks of age. To adapt them to being picked up, daily handling began immediately and lasted for several weeks. Following this acclimation period, a food-deprivation regimen was initiated during which each rat was given free access to food (lab blocks) for about 1.5 hr each day. Training in the experimental chambers began when the rats reached a weight of about 250 g (about 4 months old). Once training began, the rats were fed, in their home cages, for at least 1 hr following each daily session (usually, between 1 and 1.5 hr). With this regimen, the rats gradually gained weight, leveling off between 315 and 345 g, a weight that was sustained for the duration of the project. This weight range has been described as an appropriate and effective deprivation level for male Long Evans hooded rats (Ator, 1991).

The squads of rats differed in age and prior experimental histories. At the beginning of

this project the rats in Squad 1 were about 4 months old, those in Squad A were about 15 months old, and those in Squad B were about 7 months old. The rats in Squad A had spent the previous year in a study involving choice between different signaled delays of food reinforcement; the rats in Squad B had spent the previous 4 months responding under various multiple schedules of food reinforcement; the rats in Squad 1 had been involved in no prior projects.

Apparatus

Each rat was assigned to one of four experimental chambers for the duration of the experiment. The chambers were each 30 cm wide by 32 cm deep by 30 cm high, constructed of sheet metal (top and three sides), clear plastic (rear door, 21 cm by 30 cm), and stainless steel rods (0.7 cm diameter) spaced 1 cm apart (floor). Food pellets were delivered into a small metal food tray and were accessible via a square opening (4.4 cm by 4.4 cm) in the middle of the front panel, 4.3 cm above the floor. An electrical pulse to a Gerbrands pellet dispenser located behind the front panel caused a 45-mg Noyes Formula A pellet to drop into the food tray. The operation of the dispenser made a click, and the pellet landing in the tray made a plinking sound a fraction of a second later; both sounds were easily heard by a person standing several meters from the chamber. Pellet deliveries were not signaled by any change in illumination.

A translucent plastic key (a Lehigh Valley Electronics pigeon key) was mounted behind a round hole (1.9 cm diameter) through the left wall. The center of the key was 5.1 cm toward the rear of the chamber from the front wall and 6.2 cm above the floor. The key could be illuminated from behind with either a blinking white light (two flashes per second) or a continuously illuminated (steady) white light (two GE No. 1829 bulbs operated at 28 VDC). A response was recorded when the key was pushed a distance of about 1.5 mm (measured at the center) with a force of at least 0.18 N (three of the four chambers) or 0.3 N (the fourth chamber; Rats 14, A4, and B3). Such responses produced a brief click from a small snap-action switch connected to the key. The rats were

observed nearly always to poke the key with their noses; occasionally they used their paws.

In the back left corner of the chamber (approximately 24 cm from the front wall and 2.5 cm above the floor) was a small hole through which a metal drinking spout extended about 1 cm into the chamber. The spout was attached to a water bottle suspended outside the chamber and allowed free access to water during each session.

The four chambers were placed, two to a shelf, on a cart. The chambers were not sound shielded; it was apparent that each rat quickly learned to go to its food tray only when its own feeder operated. The four chambers typically operated at the same time, although out of phase. As an occasional check, however, an experimental session was conducted with 1 rat in isolation. These sessions provided no evidence of systematic differences attributable to whether or not the other chambers were operating.

The experimental sessions were conducted with the lights in the room darkened. A dim light, however, entered the room from the hallway through a translucent window, permitting easy observation of the rats regardless of whether the key was illuminated or darkened. The key provided the only source of illumination in the chambers (i.e., there was no houselight or feeder light). Experimental events were controlled and data recorded by a special-purpose computer connected to each chamber (Walter & Palya, 1984).

Procedure

For the rats in Squad 1, key poking was established by reinforcing successively closer approximations with food pellets (shaping). Following shaping, the rats were given a few days' training on progressively longer VI schedules and then shifted to the first condition of the experiment. The rats in Squads A and B already had considerable experience obtaining pellets by key poking; they began their first condition without any additional preliminary training.

For both conditions, the basic procedure was a two-component multiple schedule. That is, during each session the schedule alternated every 50 s between a component signaled by a blinking light (two flashes per second) and a component signaled by a steady light. Dark periods (5 s) separated each compo-

nent. Each daily session consisted of 40 blink-steady cycles (80 total components) and lasted approximately 87 min. Within a condition, one schedule of reinforcement was correlated with the blinking keylight and a different schedule of reinforcement was correlated with the steady keylight. (Responding had no programmed consequences during the 5-s blackouts between keylight components.)

For the first condition, the rich-reinforcement discriminative stimulus (steady light for the rats in Squad 1; blinking light for the rats in Squads A and B) was correlated with a VI 1-min schedule (60 pellets per hour), and the lean-reinforcement discriminative stimulus (blinking for the rats in Squad 1; steady for the rats in Squads A and B) was correlated with a VI 4-min schedule (15 pellets per hour). Technically, this procedure is a multiple VI 1-min VI 4-min schedule. Each of the VI schedules (for this and the second condition) was comprised of 16 different intervals and provided a roughly constant probability of reinforcement in time since the last reinforcer (Fleshler & Hoffman, 1962; Hantula, 1991). Upon the first presentation of a component in a session and immediately following each presentation of an assigned reinforcer, one of the 16 intervals was randomly selected (with replacement). Once an interval had assigned a reinforcer, the next interval did not begin to time until after the assigned reinforcer had been delivered. If an assigned reinforcer had not been delivered when a component ended, the reinforcer assignment was held over to the next presentation of that component. Likewise, at the end of a 50-s component, the time elapsed toward completion of an interval was saved so that the interval resumed timing at the same point when the program returned to that component.

For the second condition, during the rich-reinforcement discriminative stimulus, key pokes were reinforced by a VI 4-min schedule (15 pellets per hour) and a VT 1.33-min schedule operated concurrently, delivering pellets independently of any response (45 pellets per hour), for a total of 60 reinforcers per hour. Thus, approximately 75% of the reinforcers were presented independently of key poking. The lean-reinforcement discriminative stimulus again was correlated with a

Table 1

The order of conditions, the schedules of reinforcement that were in effect in the presence of each of the two discriminative stimuli, the type of discriminative stimulus that signaled the rich-reinforcement component during training (blinking keylight or steadily illuminated keylight), and the number of training sessions preceding each extinction test. The rich-reinforcement component of the multiple schedule provided an average of about four times as many food pellets as did the lean-reinforcement component.

Condition	Light status during rich			Sessions until extinction		
	Squad 1	Squad A	Squad B	Squad 1	Squad A	Squad B
VI 1 min (rich) VI 4 min (lean)	Steady	Blink	Blink	21	27	30
VI 4 min + VT 1.33 min (rich) VI 4 min (lean)						
Extinction 1	Steady	Blink	Blink	43	30	30
Extinction 2	Steady	Blink	Blink	25	15	12

VI 4-min schedule of reinforcement (i.e., 15 reinforcers per hour). Technically, this procedure is a multiple concurrent VI 4-min VT 1.33-min VI 4-min schedule. The VT schedules were constructed and operated as described above for the VI schedules, except that the food pellet was delivered immediately after an interval of the VT elapsed, regardless of responding.

Tests for resistance to extinction were carried out exactly like baseline training for the relevant condition except that after the 10th cycle (i.e., during the last 30 cycles), no food pellets were delivered. (Note that the first 10 cycles continued baseline training.) One extinction session was given following training under the first condition (all reinforcers were response dependent). Two extinction sessions were given following training under the second condition; they were separated by additional sessions of baseline training.

With few exceptions, sessions were conducted 7 days per week. Table 1 summarizes the training conditions and lists the number of sessions given prior to each extinction session. The number of training sessions appeared to be more than sufficient to generate stable baseline rates of responding in the presence of the two discriminative stimuli (i.e., no apparent upward or downward trends in response rates for any of the rats). Some results of the baseline training preceding the first two extinction sessions were presented, in a different context, in Shull *et al.* (2001) for the rats in Squads A and B.

Interresponse times were recorded separately for each discriminative stimulus during the last three baseline sessions of each con-

dition (but only during the last session prior to the third extinction test for the rats in Squads A and B). The IRTs were also recorded during extinction sessions. The computer listed IRTs in units of 0.01 s, but our system actually registered IRTs to about the nearest 0.1 s. The IRT for the first response in each 50-s discriminative stimulus component was measured from the start of the component; the IRTs of all subsequent responses during the component were measured from the previous response.

RESULTS

Baseline Performance and Resistance to Extinction

Figure 1 shows, for the rats in Squad 1, how response rate in the presence of the two discriminative stimuli changed over consecutive 10-cycle segments of the extinction sessions. Because the y axis is scaled logarithmically, vertical distance indicates relative change. Thus, the shallower the slope of a plot, the more resistant the indicated response rate was to extinction. Differential persistence is indicated by plots that diverge. As expected from prior results, response rate declined more slowly during the extinction sessions in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. This was true regardless of whether all or only about 25% of the reinforcers in the presence of the rich-reinforcement signal had been response dependent (compare the plots in the left column of panels of Figure 1 with those in the middle and right panels).

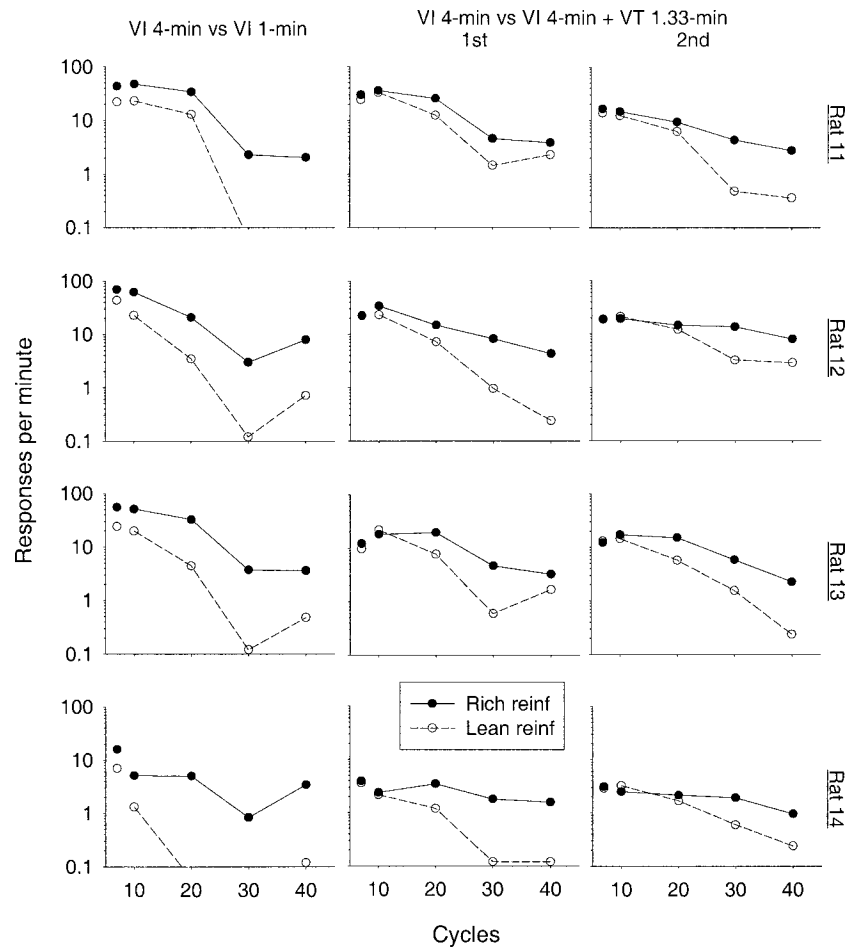


Fig. 1. Response rates during extinction sessions for each of the 4 rats in Squad 1. The schedules of reinforcement that had been in effect during baseline training in the presence of the two discriminative stimuli of the multiple-schedule pair are indicated at the top. Note that for the panels in the left column all the reinforcers had been response dependent during the rich-reinforcement discriminative stimulus. For the middle (first test) and right (second test) columns, about 75% of the reinforcers during the rich-reinforcement discriminative stimulus had been delivered independent of responding during baseline (i.e., on a VT schedule). The unconnected data points at the far left show the mean response rates over the last five sessions of baseline training. Each connected point shows the response rate averaged over the prior 10 cycles during the extinction session. Because the baseline training schedules were in effect during the first 10 cycles of the extinction test session, the first connected point in each plot provides another estimate of baseline response rate. Because the y axis is scaled logarithmically, the slopes of the lines indicate relative change in response rate. Plots that drop below 0.1 responses per minute indicate that no responses occurred during that block of cycles; a value of 0.06 was arbitrarily entered for those no-response occasions.

Figure 1 also shows the response rates in the presence of the two discriminative stimuli at the end of baseline training (unconnected points at the far left in each panel). Moreover, because the baseline training schedules operated during the first 10 cycles of the extinction session, the initial point in each plot serves as an additional sample of baseline response rate. As expected, rates of key poking in the presence of the two discriminative

stimuli differed relatively more from each other when all the reinforcers in the presence of the rich-reinforcement signal were response dependent than when only 25% were (compare the vertical distances between the pairs of baseline points in the left column of panels of Figure 1 with those in the middle and right panels). In short, the data in Figure 1 replicate (with rats and key poking) the well-established finding that resistance to ex-

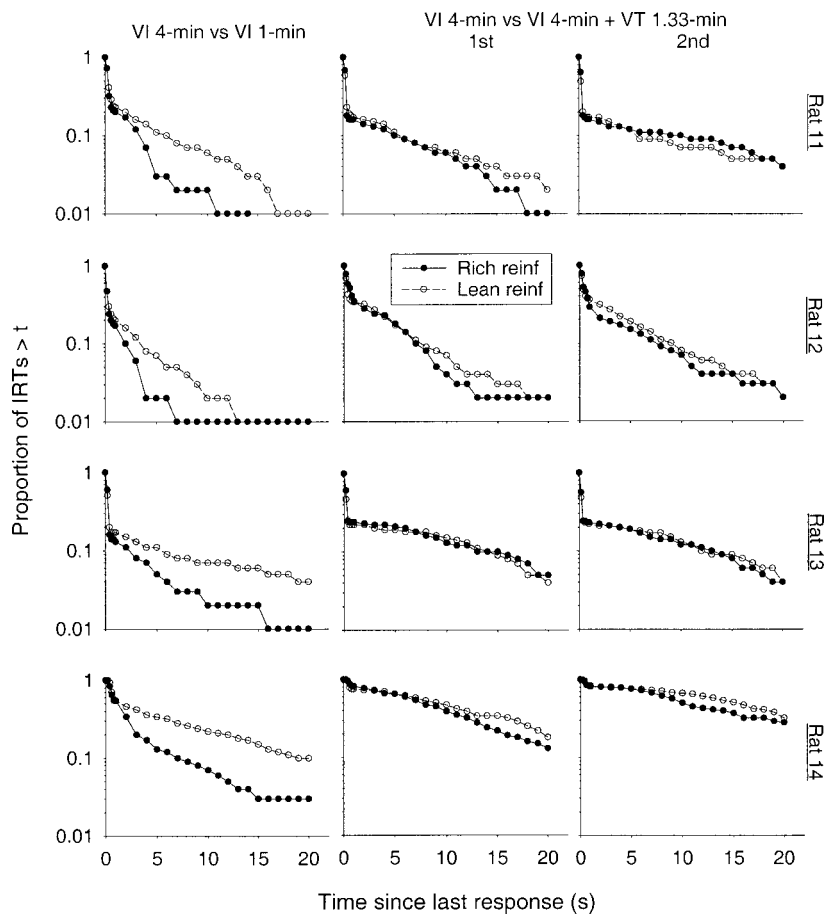


Fig. 2. Log survivor plots of IRTs (semilog plots of the proportion of IRTs longer than t s) from the baseline sessions that immediately preceded the corresponding extinction sessions represented in Figure 1. The schedules of reinforcement that were in effect during baseline training in the presence of the two discriminative stimuli of the multiple-schedule pair are indicated at the top.

tion in the presence of a discriminative stimulus is related to the relative rate of reinforcement that has occurred in the presence of the stimulus but not to the relative reinforcement of the response or to the relative rate of the response between the two discriminative stimuli during baseline.

We next consider IRT data in order to determine the bout-initiation rates and the number of responses per bout during baseline training. For this purpose, log survivor plots have proven to be useful (see Shull *et al.*, 2001, for an extensive discussion). These are semilogarithmic plots of the proportion of IRTs longer than t time units (here, seconds). Figure 2 shows log survivor plots of IRTs from the baseline training session that occurred immediately prior to each of the ex-

tinguishing sessions that produced the data shown in Figure 1. Consider first the plots in the left column (VI 1 min and VI 4 min). Most of the plots have what is sometimes called a *broken-stick* appearance (e.g., Clifton, 1987), which indicates that responses have been emitted at two different rates, as implied by the pause-bout conception of response rate (Howard, 1963, pp. 14–16; Sibley, Nott, & Fletcher, 1990; Tolkamp, Allcroft, Austin, Nielsen, & Kyriazakis, 1998). The two plots in each panel appear to break at about the same point along the y axis, but then the limb for the lean-reinforcement schedule (VI 4 min) remains flatter than the limb for the rich-reinforcement schedule (VI 1 min). As discussed by Shull *et al.*, if the limb to the right of the break in each plot is reasonably

linear (with logarithmic y axis), its slope can be interpreted as indicating the bout-initiation rate. The steeper the slope, the higher the bout-initiation rate; thus, for each of the rats, the richer VI schedule generated the higher bout-initiation rate (Figure 1, left column). If a bout-initiation limb is extrapolated back to where it would intersect the y axis, that point indicates the proportion of total responses that are bout initiations, the inverse of which is the mean number of responses per bout. Thus, the longer the initial limb is before the break, the larger the number of responses per bout (i.e., bout length). It is apparent, for example, that bout lengths were longer for Rat 13 than for Rat 14. But for each rat, the bout lengths were similar for the two schedules of the pair. Considering the panels in the left column, then, the value of the VI schedule had a systematic effect on the bout-initiation rate (i.e., the diverging limbs to the right of the break) but little effect on the number of responses per bout (as implied by the similar y -axis intercepts of the two bout-initiation limbs in each panel).

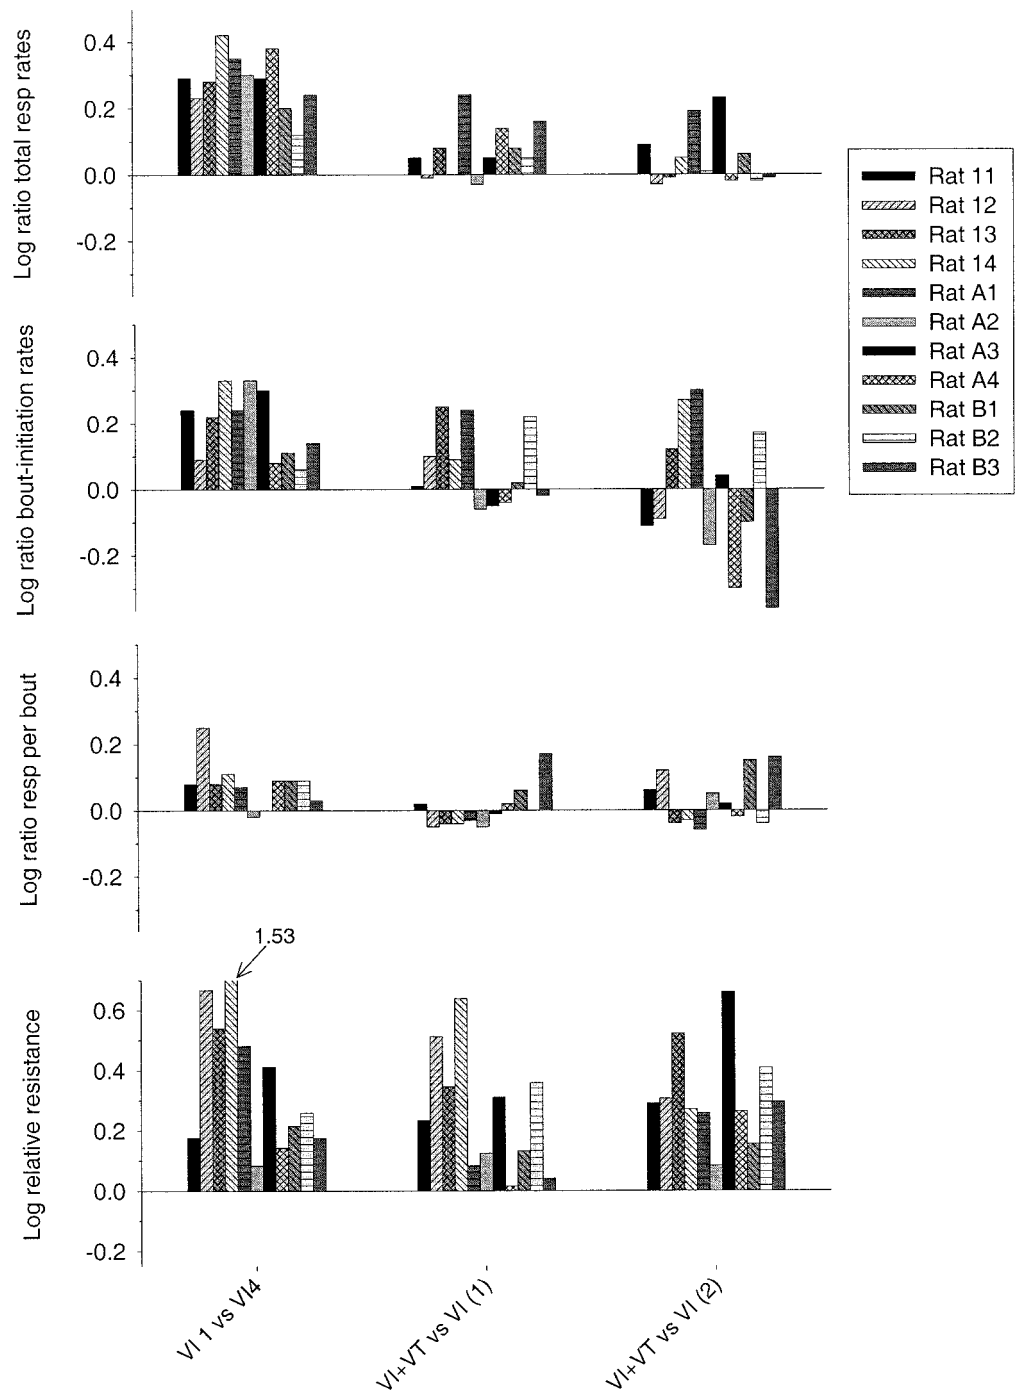
When 75% of the reinforcers in the rich-reinforcement component were response independent (middle and right columns of Figure 2), the bout-initiation rates were similar in the presence of the two discriminative stimuli. Indeed, there is no systematic difference in any feature of the log survivor plots of IRTs between the VI (lean) and the VI + VT (rich) schedules. Thus, the differential resistances to extinction (Figure 1) were not correlated with differences during baseline in total response rate (as has already been well established) or with differences in two components of total response rate—namely, bout-initiation rate and number of responses per bout.

The detailed data shown in Figures 1 and 2 are from the 4 rats in Squad 1. Instead of presenting the corresponding detailed data from the remaining 7 rats, we provide summary measures of the relevant aspects of performance for all 11 rats in Figure 3. The data in the top three rows indicate performance during baseline. The performance measures are summarized as logarithms of ratios; a value of zero indicates no difference in the performance measure between the two discriminative stimuli [i.e., $\log(1) = 0$]. Positive values indicate that the performance measure

was higher for the rich-reinforcement signal than for the lean-reinforcement signal. Negative values indicate the reverse. The performance measure in the bottom row indicates the relative resistance to extinction, again in logarithmic form, so that positive values indicate greater resistance in favor of the rich-reinforcement signal.

More specifically, total response rates (top row) were calculated in the usual way: Total key pokes in the presence of a discriminative stimulus were divided by the time spent in the presence of the discriminative stimulus. Bout-initiation rates (second row) and the mean numbers of responses per bout (third row) were calculated from log survivor plots by the method described by Shull et al. (2001). A straight line (log y axis) was fit to the segment of the log survivor plot from 1 s through 12 s. The slope of that best fitting line was then taken as an estimate of the bout-initiation rate; the inverse of the y -axis intercept gave the mean number of responses per bout. The log relative resistance index (cf. Grace & Nevin, 1997, 2000) was calculated as $\log(R_x/R_0) - \log(L_x/L_0)$, where R_x indicates the total response rate in the presence of the rich-reinforcement signal during the true extinction portion of the extinction session (i.e., over Cycles 11 through 40), and R_0 indicates the total response rate in the presence of the rich-reinforcement signal during baseline; L_x and L_0 indicate the corresponding response rates in the presence of the lean-reinforcement signal. The ratios (R_x/R_0 and L_x/L_0) express response rate in extinction as a proportion of baseline response rate. Thus, the difference between the logarithms of these ratios indicates whether response rate was more resistant to extinction, relative to its baseline level, in the presence of the rich-reinforcement signal (positive values) or in the presence of the lean-reinforcement signal (negative values). The size of the difference between the logarithms indicates the magnitude of the relative difference between the resistances for the two discriminative stimuli.

The purpose of Figure 3 is to facilitate inspection for any general relation between performance during baseline and resistance to extinction. No such relation is apparent. The relative-resistance-to-extinction indexes (bottom row) were positive in every case, and their magnitudes did not differ systematically



as a function of whether all or only about 25% of the reinforcers in the presence of the rich-reinforcement discriminative stimulus had been response dependent. In contrast, the ratio of total response rates (top row) and the ratio of bout-initiation rates (second row) usually decreased toward equality when the percentage of response-dependent reinforcers in the presence of the rich-reinforcement discriminative stimulus went from 100% (left clusters) to 25% (middle and right clusters). Indeed, in some cases in which only 25% of the reinforcers were response dependent, the response rates and bout-initiation rates were relatively higher in the presence of the lean-reinforcement signal (negative values in the top two rows, middle and right clusters). The baseline average number of responses per bout (third row from the top) tended to be somewhat higher in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal when all the reinforcers were response dependent, but the relative difference was small and not entirely consistent across rats. There was no systematic difference between the discriminative stimuli in the number of responses per bout when only 25% of the reinforcers in the presence of the rich-reinforcement signal were response dependent. In all, these patterns of results are entirely consistent with the patterns discussed earlier in relation to the more detailed presentations in Figures 1 and 2 for the rats in Squad 1. They confirm that relative resistance to extinction is not related to baseline differences between the two discriminative stimuli in any of the performance measures that we considered. Instead, relative resistance to extinction is directly related to the relative reinforcement that has been cor-

related with the discriminative stimuli during baseline training.

Bout-Initiation Rate and Bout Length During Extinction

One of our aims in obtaining IRTs was to track the changes in bout-initiation rate and responses per bout throughout extinction sessions. We did this two different ways. First, we attempted to measure bout-initiation rate and responses per bout from log survivor plots of IRTs recorded during different segments of each extinction session. This approach was not particularly successful, however, mainly because the number of IRTs per sample was often quite small, especially for segments late in extinction. When the sample size is small, the estimates from the limbs of the log survivor plots tend to be unreliable. Although there is no criterion number for a minimum sample size, it may be worth noting that the baseline log survivor plots in Figure 2 were based on a mean sample size of 695 IRTs. In contrast, many of the 10-cycle samples from extinction sessions contained fewer than 15 IRTs. Nevertheless, we describe some results of the within-extinction analysis based on log survivor plots because the results seemed at least suggestive of interesting conclusions.

Our second approach to estimating bout-initiation rate and responses per bout within extinction sessions followed what might be called the *cutoff IRT* method (e.g., Mellgren & Elsmore, 1991). Initially, some IRT duration (e.g., 1 s) is selected as a criterion value. All responses that end an IRT that is longer than the criterion are classed as bout initiations; all other responses are classed as within-bout responses. Considering where the

←

Fig. 3. Summary of performance measures from baseline (top three rows) and from the corresponding extinction session (bottom row) for all 11 rats (Squads 1, A, and B). Each bar represents the data point from an individual rat (ordered from left to right in each cluster). The clusters correspond to one of the three resistance-to-extinction assessments, as indicated by the labels at the bottom. For each training condition, the rate of reinforcement during the rich-reinforcement signal was approximately four times that in the presence of the lean-reinforcement signal (approximately 60 vs. 15 reinforcers per hour). Total response rate in the presence of each of the two discriminative stimuli was computed from the session totals. Bout-initiation rate and average number of responses per bout were estimated from fits to the limbs of log survivor plots, as described in the text. The bars in the top three rows show the common logarithm of the ratios of these response measures (rich reinforcement to lean reinforcement). The measure of relative resistance to extinction is described in the text. For all four rows of data, positive values indicate that the performance measure was higher in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. Negative values indicate the reverse. The performance measures that were the basis of these ratios are listed in the Appendix.

breaks occurred in the baseline log survivor plots (e.g., Figure 2), we set the criterion duration at 1 s. This value seemed sufficiently to the right of the break to exclude most of the within-bout responses from the bout-initiation class and sufficiently close to the break so as not to misclassify too many bout initiations as within-bout responses. (In fact, we repeated the analysis with cutoff IRT values set at 0.7 s, 5 s, and 10 s; the pattern of results to be described shortly was essentially the same regardless of the cutoff duration.) There are some disadvantages with the cutoff IRT method of estimating bout-initiation rate and responses per bout, as discussed by Shull *et al.* (2001), but none are seriously problematic for the present purpose of tracking relative trends during extinction. One advantage of the cutoff IRT method is that it can be applied even with small sample sizes. One simply counts the number of responses of a given type and divides by the time base to determine a rate. (If bout-initiation rates are high or bout lengths long, the time during bouts should be subtracted from the time base for calculating bout-initiation rate. That correction would have had a negligible effect in the present data, and we did not make that correction.)

Log survivor estimates. Recall that the first 10 cycles of the 40-cycle extinction session were a continuation of baseline training. We used the IRTs from those 10 cycles to estimate baseline performance measures. Then we used the IRTs from Cycles 11 through 20 to estimate performance measures early in extinction and those from Cycles 21 through 40 to estimate performance measures during the latter part of the extinction session. Note that these samples include a larger block of cycles as extinction progressed. This was done to generate a larger sample late in extinction than would have been possible had the samples been derived from consecutive 10-cycle blocks.

Figure 4 shows log survivor plots from the segments of the extinction session following training with VI 1 min and VI 4 min for the 4 rats in Squad 1. Although it is difficult to discern a consistent pattern, there was some tendency for the bout-initiation rates to decrease over the course of extinction. Such a pattern is most evident for Rat 12 (left column): The tails of the log survivor plots tend

to rotate upward as extinction progressed. Also, there was some tendency for the bout-initiation rate to approach zero (either a horizontal limb or no plot due to a lack of responding during the segment) more quickly in the presence of the lean-reinforcement signal (right column) than in the presence of the rich-reinforcement signal (left column). There was also some evidence that the number of responses per bout decreased as extinction progressed. Such an effect is shown by an upward movement of the y-axis intercept, extrapolating the limb back to the y axis, as extinction continued. One notable feature of these log survivor plots is that they retain their broken-stick appearance, even late in extinction. That is, responding remains pause-bout in character throughout much of extinction.

As can be inferred from the variability apparent in the sample of log survivor plots in Figure 4, the numerical estimates of bout-initiation rate and average number of responses per bout that were derived from such log survivor plots produced extinction functions that were highly variable among the 11 rats. Plots of the mean values over the 11 rats, however, provide some evidence of orderly trends, as shown in Figure 5. Mean bout-initiation rate decreased over the extinction session and did so differentially, declining more slowly in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. Moreover, this differential persistence in bout-initiation rate was apparent whether the reinforcers in the presence of the rich-reinforcement signal had been 100% response dependent (left panel) or only about 25% response dependent (middle and right panels). The average number of responses per bout also declined over the extinction session, to a slightly greater extent in the presence of the lean-reinforcement signal than in the presence of the rich-reinforcement signal. The differential persistence of average number of responses per bout, however, was considerably smaller than that for bout-initiation rate. (The conclusions from the plots in Figure 5 would have been similar had median values been plotted instead of means.)

Cutoff IRT method. Figure 6 likewise shows that the rate of bout initiations (defined by the cutoff IRT method) declined more slowly




Fig. 4. Log survivor plots of IRTs from different blocks of cycles during the extinction test session following baseline training under a multiple VI 1-min VI 4-min schedule of reinforcement. Each row of plots shows the data from a different rat in Squad 1. The plots in the left column show data from the rich-reinforcement discriminative stimulus; those in the right column show data from the lean-reinforcement discriminative stimulus. Cycles 1 through 10 provided reinforcement as under baseline; extinction was in effect during the remaining 30 cycles. Missing plots indicate that no responses occurred in the presence of the discriminative stimulus during that block of cycles.

throughout extinction in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. The size of the differential persistence is about the same, regardless of whether all or only 25% of the reinforcers in the presence of the rich-reinforcement signal had been response dependent. Here the differential persistence is apparent in the plots for the individual rats (Squad 1).

An index of differential persistence, similar to the one shown in the bottom row of Figure

3, can be computed for the extinction plots of bout-initiation rate, such as those in Figure 6 (cf. Grimes & Shull, 2001). For this index, the bout-initiation rates during the first 10 cycles of the extinction session were taken as baseline values. Then for each discriminative stimulus, the bout-initiation rate in each consecutive 10-cycle segment of extinction was expressed as a proportion of its baseline rate, and the logarithm of the proportion was determined. The resistance score for each discriminative stimulus is the mean of these log-




Fig. 5. Within-session extinction functions for bout-initiation rate (top row) and average number of responses per bout (bottom row) derived from log survivor plots. The points show geometric mean values computed over the 11 rats. The first 10 cycles of the extinction test session provided reinforcers as during baseline. The schedules in effect during the immediately prior baseline training are indicated at the top. When the bout-initiation rate for a block of cycles was zero (or undefined because of no responses), a value of 0.04 per minute was entered. This provided a conservative estimate. Such a correction was required for 14 of the 132 entries in the data set (i.e., for 7 of 44, 2 of 44, and 5 of 44, respectively, for the left, middle, and right panels). Note that the cycle blocks, from left to right in each panel, are partly overlapping and progressively larger.

arithms. The relative resistance index is the difference between these mean logarithms for the two discriminative stimuli. Formally, this index is computed as $(1/3) \sum \log (R_x/R_0) - (1/3) \sum \log (L_x/L_0)$, where the terms are as described earlier in conjunction with Figure 4. For comparison, the same kind of index was computed for the extinction plots for total response rate, as shown in Figure 1. Table 2 lists these mean log relative resistance indexes for all 11 rats, for all three extinction tests, and for both bout-initiation rate and total response rate. For bout-initiation rate, 32 of the 33 indexes are positive in sign, indicating greater resistance to extinction in the presence of the rich-reinforcement signal. For total response rate, 31 of the 33 of the indexes are positive in sign. Thus, conclusions about which discriminative stimulus engenders the more persistent responding are the same whether they are based on total response rate or on bout-initiation rate.

Further evidence of the consistency between total response rate and bout-initiation rate in their differential resistances to extinction is shown in Figure 7. The mean log relative resistance indexes based on the two performance measures correspond closely. For the first, second, and third extinction tests, the r^2 values for the best fitting straight lines are .76, .93, and .84, respectively. In other words, for a particular rat and extinction test, if differential resistance was relatively weak (or strong) based on total response rate, it was also likely to be weak (or strong) based on bout-initiation rate. This was true whether all or only 25% of the reinforcers during the rich-reinforcement signal were response dependent.

Figure 8 shows that the average number of responses per bout (derived from the cutoff IRT method) tended to decrease over the course of extinction for 3 of the 4 rats in Squad 1. The number of responses per bout

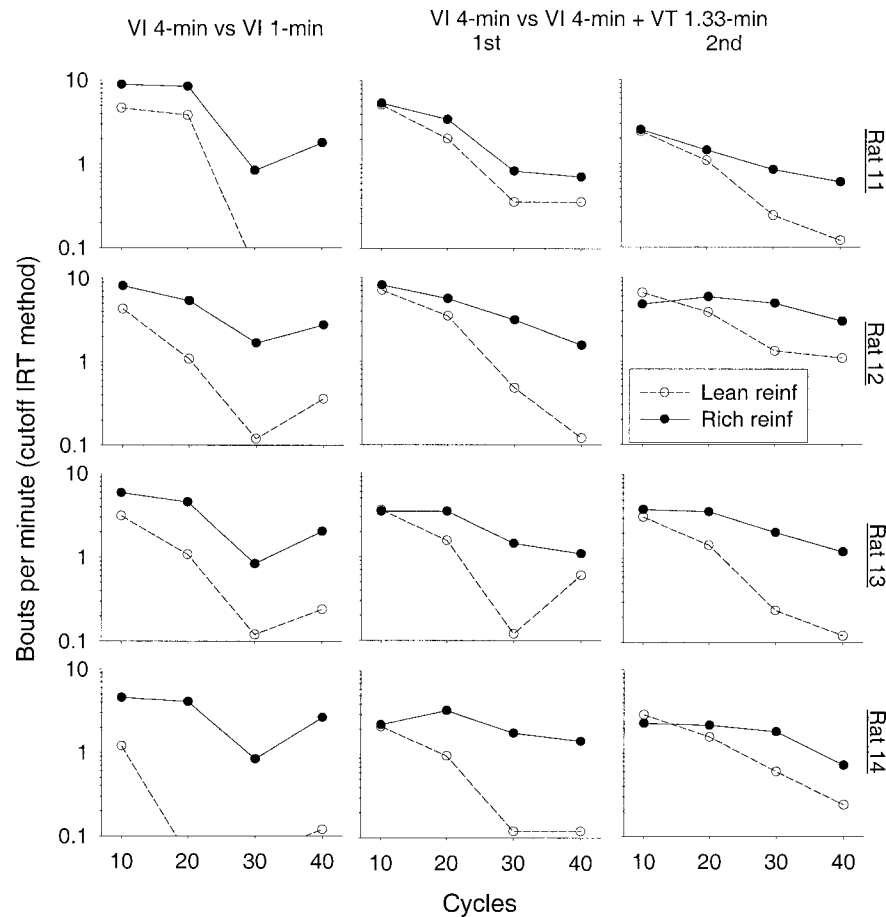


Fig. 6. Within-session extinction functions for bout-initiation rate (cutoff IRT method) for the rats in Squad 1. A response that followed an IRT that was equal to or longer than 1 s was classed as a bout initiation; a response that followed an IRT shorter than 1 s was classed as a within-bout response. Note that the y axes are scaled logarithmically so that the slope indicates relative change. The schedules of reinforcement that had been in effect during baseline training in the presence of the two discriminative stimuli of the multiple-schedule pair are indicated at the top. Each point shows the response rate averaged over 10 cycles during the extinction session. Because the baseline training schedules were in effect during the first 10 cycles of the extinction test session, the first connected point in each plot provides an estimate of baseline bout-initiation rate. Plots that drop below 0.1 responses per minute indicate that no responses occurred during that block of cycles; a value of 0.06 was arbitrarily entered for those no-response occasions.

for the 4th rat (Rat 14) was already close to its minimum value of 1 under baseline, so there was virtually no room for further decline. For a few extinction tests, the number of responses per bout were more persistent in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal (e.g., Rat 13, left column). But overall, there was little evidence of any consistent differential resistance to extinction in the number of responses per bout. The extinction plots shown in Figure 8 for the rats in Squad 1 were similar to those for the other

rats (not shown). The trends are also consistent with those apparent in Figure 5 (bottom row) for the group-mean responses per bout derived from the log survivor method. That is, bout length tended to decline over the course of extinction, and that decline was at most only slightly differential between the two discriminative stimuli.

It is possible to compute a mean log relative resistance index for bout length, analogous to the indexes for total response rate and bout-initiation rate shown in Table 2. Indeed, it can be shown that such an index (for

Table 2

Mean log relative-resistance-to-extinction indexes for each rat and for each of the three extinction sessions. The training schedules are indicated at the top. For each extinction test, the index is shown based on total response rate and bout-initiation rate. Positive values of the index indicate greater resistance to extinction in the presence of the rich-reinforcement signal; negative values indicate greater resistance to extinction in the presence of the lean-reinforcement signal. The calculation of the index is described in the text.

Rat	VI 1 min (rich) vs. VI 4 min (lean)		VI 4 min + VT 1.33 min (rich) vs. VI 4 min (lean)			
			Extinction 1		Extinction 2	
	Total	Bout	Total	Bout	Total	Bout
11	0.86	0.71	0.31	0.28	0.59	0.43
12	0.63	0.63	0.67	0.65	0.43	0.54
13	0.66	0.52	0.61	0.58	0.59	0.69
14	0.92	0.86	0.87	0.91	0.51	0.47
A1	0.31	0.24	0.38	0.36	0.56	0.42
A2	0.34	0.29	-0.11	0.05	0.23	0.14
A3	0.67	0.45	0.35	0.25	0.19	0.17
A4	0.60	0.55	0.24	0.26	0.25	0.25
B1	0.40	0.22	0.40	0.32	-0.02	-0.05
B2	0.78	0.41	0.66	0.53	0.56	0.48
B3	0.33	0.30	0.03	0.02	0.45	0.38




Fig. 7. The correspondence between mean log relative-resistance-to-extinction values based on total response rate with those based on bout-initiation rate (cutoff IRT method). The details of these calculations are given in the text. Each data point indicates the coordinate values for 1 of the 11 rats. Positive values indicate that the relative resistance to extinction was greater in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. As indicated in the legend, the three assessments of resistance to extinction are differentiated by the type of symbol. The data are from Table 2. The dotted diagonal line indicates perfect correspondence.

classification based on the IRT cutoff method) is simply the difference between the index for total response rate and that for bout-initiation rate. As can be determined from the data in Table 2, this difference is, on average, positive in sign but relatively small in magnitude. In other words, the differential resistance of bout-initiation rate and of bout length both contribute to the differential resistance of total response rate. But, again, the contribution from the bout-length measure is considerably smaller and less reliable than that from the bout-initiation-rate measure.

DISCUSSION

The present results replicate, with rats and key poking, an important and well-established finding about the resistance of operant behavior to extinction in the presence of two different discriminative stimuli. The discriminative stimulus that has been correlated with the higher rate of reinforcement during training will engender the more persistent responding during extinction. Moreover, this relation holds whether all or only some of the reinforcers that have occurred in the presence of the discriminative stimulus have been response dependent. Such results (e.g., Dube *et al.*, 2000; Grimes & Shull, 2001; Mace *et al.*, 1990; Mauro & Mace, 1996; Nevin *et al.*, 1990; Figures 1 and 3 in the present paper)




Fig. 8. Within-session extinction functions for responses per bout (cutoff IRT method) for the rats in Squad 1. Missing points indicate that no response occurred during that block of cycles. Other details are as given for Figure 6.

have supported the principle that resistance to change depends much more on the stimulus–reinforcer contingency than on the response–reinforcer contingency (cf. Nevin, 1992; Nevin & Grace, 2000).

The IRT analyses in the present study were undertaken in response to concerns that this principle might lack generality because it is based on total response rate as the measure of persistence. These concerns stemmed from evidence that total response rate is best viewed as a composite measure rather than a unitary one. As discussed in the introduction, it seemed possible that different conclusions might be reached about the determinants of resistance to extinction depending on whether resistance is measured by total response rate or by bout-initiation rate. For example, it seemed possible that response-independent deliveries of the reinforcer might affect total response rate by generating relatively short

bouts that occur at a relatively high rate. If so, the resistance of a response to extinction might correlate at least as well with the baseline bout-initiation rate as with the relative reinforcement in the presence of a discriminative stimulus. Alternatively, it seemed possible that the persistence of bout-initiation rate, unlike total response rate, might be negatively related to the rate of reinforcement in the presence of the discriminative stimulus. If so, that would create an ambiguity about which relation better describes a principle applicable to operant behavior in general.

It turned out, however, that there was no evidence of such complicating patterns at the level of bout-initiation rate and responses per bout. Indeed, whether the focus was on performance during baseline training (e.g., Figures 2 and 3) or during the extinction sessions (e.g., Figures 1, 5, and 6 and Table 2), conclusions about differential resistance to

extinction based on bout-initiation rate were virtually the same as those based on total response rate. With either measure of responding, resistance to extinction was greater in the presence of the discriminative stimulus that had been correlated with the higher rate of reinforcers, regardless of whether all or only 25% of those reinforcers were response dependent and regardless of the relative rate of responding (either bout initiation or total responses) during baseline.

Moreover, there was no evidence that the different baseline training schedules generated different bout lengths between the two discriminative stimuli (Figures 2 and 3). Nor was there any evidence that the different resistances of total response rate to extinction could be attributed to different susceptibilities of bouts to disruption by extinction (Figures 5 and 8). If there was any such differential susceptibility, it was small in magnitude relative to the differential persistence of bout-initiation rate. In short, the present data should alleviate concerns that conclusions about the determinants of differential resistance to extinction are limited in generality because they are based on total-response-rate measures.

We had hoped to trace the course of bout-initiation rate and number of responses per bout during extinction from estimates derived from log survivor plots, but we were unable to do so effectively for individual rats because of the small sample sizes. The group-mean data (Figure 5), however, indicated that bout-initiation rate was more persistent in the presence of the rich-reinforcement signal than in the presence of the lean-reinforcement signal. A similar pattern of differential resistance to extinction was apparent in the data from the individual rats in which bout-initiation rate was determined by the cutoff IRT method (Figures 5 and 6 and Table 2). The consistency between the two methods of measuring bout-initiation rate may encourage confidence in the reliability of the pattern of results despite the fact that one of the methods yielded orderly extinction functions only at the level of group means.

Although the resistance of bout initiations to extinction may sometimes be negatively related to the rate of reinforcement during training (Mellgren & Elmsore, 1991), it ap-

pears to be positively related when the rates of reinforcement are correlated with different discriminative stimuli that alternate during the training sessions (see also Elmsore, 1971, 1979). Such results are consistent with the results and interpretation offered by Cohen (1998), namely that the discriminative stimuli (and their correlated rates of reinforcement) must alternate fairly frequently in order for the stimulus-reinforcer contingencies to affect persistence.

Although the conclusions regarding resistance to extinction are similar whether they are based on changes in total response rate or bout-initiation rate, the bout-initiation-rate findings facilitate the interpretation of the effects of extinction. For example, previous research has demonstrated that bout-initiation rate is particularly sensitive to relative reinforcement variables (Shull *et al.*, 2001). Thus, the decline in bout-initiation rate during extinction might be interpreted as reflecting a decline in relative reinforcement value during extinction. That the bout-initiation rate declined differentially in the presence of the two discriminative stimuli suggests that the relative reinforcement value was declining at different rates. Such a conclusion seems to be entirely consistent with the theoretical interpretation of behavioral momentum offered by Grace and Nevin (1997; see also Nevin & Grace, 2000), which holds that preference and resistance to change are functions of a common effect of reinforcement.

The smaller differential decline in responses per bout during extinction is consistent with the pause-bout conception of responding. Bout initiations can be conceptualized as the first member of a chain (analogous to a changeover in concurrent schedules), and the bout itself can be conceptualized as the second (i.e., terminal) member of the chain. A well-documented result is that initial members of chains are less resistant to change than are terminal members (Mellon & Shull, 1986; Nevin, Mandell, & Yarensky, 1981).

It has been known for a long time that long IRTs change more than short IRTs during extinction (Blough, 1963; Schaub, 1967). Under the pause-bout interpretation, responses that end long IRTs will tend to be bout initiations, and many of the responses that end short IRTs are likely to be within-bout responses. Thus, our finding that bout-initia-

tion rate is particularly sensitive to extinction is consistent with the earlier findings based on IRT duration.

In conclusion, the data for both total response rate and bout-initiation rate support a key feature of behavioral momentum theory (Nevin, 1992; Nevin & Grace, 2000), namely that the rate of operant behavior under constant conditions and the resistance of that behavior to change are products of different types of contingency. Response rate under constant conditions depends heavily on the stimulus–response–reinforcer (operant) contingency, whereas resistance to change depends heavily on the stimulus–reinforcer (Pavlovian) contingency.

REFERENCES

- Ator, N. A. (1991). Subjects and instrumentation. In I. H. Iversen & K. A. Lattal (Eds.), *Techniques in the behavioral and neural sciences: Vol. 6. Experimental analysis of behavior* (Part 1, pp. 1–62). Amsterdam: Elsevier.
- Blough, D. S. (1963). Interresponse time as a function of continuous variables: A new method and some data. *Journal of the Experimental Analysis of Behavior*, 6, 237–246.
- Branch, M. N. (2000). Gaining (on) momentum. *Behavioral and Brain Sciences*, 23, 92–93.
- Burgess, I. S., & Wearden, J. H. (1986). Superimposition of response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, 45, 75–82.
- Clifton, P. G. (1987). Analysis of feeding and drinking patterns. In F. M. Toates & N. E. Rowland (Eds.), *Feeding and drinking* (pp. 19–35). Amsterdam: Elsevier.
- Cohen, S. L. (1998). Behavioral momentum: The effects of the temporal separation of rates of reinforcement. *Journal of the Experimental Analysis of Behavior*, 69, 29–47.
- Davison, M., & Charman, L. (1986). On the measurement of time allocation on multiple variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 46, 353–362.
- Doughty, A. H., & Lattal, K. A. (2001). Resistance to change of operant variation and repetition. *Journal of the Experimental Analysis of Behavior*, 76, 195–215.
- Dube, W. V., Mazzitelli, K., Lombard, K. M., & McIlvane, W. J. (2000). Assessing behavioral momentum in humans with mental retardation and unstable baselines. *Experimental Analysis of Human Behavior Bulletin*, 18, 6–11.
- Elsmore, T. F. (1971). Effects of response effort on discrimination performance. *Psychological Record*, 21, 17–24.
- Elsmore, T. F. (1979). Supplementary feeding effects on multiple fixed-ratio performance of rhesus monkeys. *Bulletin of the Psychonomic Society*, 13, 367–370.
- Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 5, 529–530.
- Galbicka, G., & Kessel, R. (2000). Experimenter momentum and the effect of laws. *Behavioral and Brain Sciences*, 23, 97–98.
- Gilbert, T. F. (1958). Fundamental dimensional properties of the operant. *Psychological Review*, 65, 272–282.
- Grace, R. C., & Nevin, J. A. (1997). On the relation between preference and resistance to change. *Journal of the Experimental Analysis of Behavior*, 67, 43–65.
- Grace, R. C., & Nevin, J. A. (2000). Comparing preference and resistance to change in constant- and variable-duration schedule components. *Journal of the Experimental Analysis of Behavior*, 74, 165–188.
- Grimes, J. A., & Shull, R. L. (2001). Response-independent milk delivery enhances persistence of pellet-reinforced lever pressing by rats. *Journal of the Experimental Analysis of Behavior*, 76, 179–194.
- Hantula, D. A. (1991). A simple BASIC program to generate values for variable-interval schedules of reinforcement. *Journal of Applied Behavior Analysis*, 24, 799–801.
- Howard, R. A. (1963). *Nuclear physics*. Belmont, CA: Wadsworth.
- Mace, F. C., Lalli, J. S., Shea, M. C., Lalli, E. P., West, B. J., Roberts, M., et al. (1990). The momentum of human behavior in a natural setting. *Journal of the Experimental Analysis of Behavior*, 54, 163–172.
- Mauro, B. C., & Mace, F. C. (1996). Differences in the effect of Pavlovian contingencies upon behavioral momentum using auditory versus visual stimuli. *Journal of the Experimental Analysis of Behavior*, 65, 389–399.
- Mechner, F. (1992). *The revealed operant: A way to study the characteristics of individual occurrences of operant responses*. Cambridge, MA: Cambridge Center for Behavioral Studies.
- Mellgren, R. L., & Elsmore, T. F. (1991). Extinction of operant behavior: An analysis based of foraging considerations. *Animal Learning & Behavior*, 19, 317–325.
- Mellon, R. C., & Shull, R. L. (1986). Resistance to change produced by access to fixed-delay versus variable-delay terminal links. *Journal of the Experimental Analysis of Behavior*, 46, 79–92.
- Neuringer, A., Kornell, N., & Olufs, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 79–94.
- Nevin, J. A. (1992). An integrative model for the study of behavioral momentum. *Journal of the Experimental Analysis of Behavior*, 57, 301–316.
- Nevin, J. A., & Baum, W. M. (1980). Feedback functions for variable-interval reinforcement. *Journal of the Experimental Analysis of Behavior*, 34, 207–217.
- Nevin, J. A., & Grace, R. C. (2000). Behavioral momentum. *Behavioral and Brain Sciences*, 23, 73–130. (includes commentary)
- Nevin, J. A., Mandell, C., & Yarensky, P. (1981). Response rate and resistance to change in chained schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 278–294.
- Nevin, J. A., Tota, M. E., Torquato, R. D., & Shull, R. L. (1990). Alternative reinforcement increases resistance to change: Pavlovian or operant contingencies? *Journal of the Experimental Analysis of Behavior*, 53, 359–379.
- Pear, J. J., & Rector, B. L. (1979). Constituents of response rate. *Journal of the Experimental Analysis of Behavior*, 32, 341–362.
- Premack, D. (1965). Reinforcement theory. In D. Levine

- (Ed.), *Nebraska Symposium on Motivation* (pp. 123–180). Lincoln: University of Nebraska Press.
- Rachlin, H., & Baum, W. M. (1972). Effects of alternative reinforcement: Does the source matter? *Journal of the Experimental Analysis of Behavior*, 18, 231–241.
- Schaub, R. E. (1967). Analysis of interresponse times with small class intervals. *Psychological Record*, 17, 81–89.
- Shimp, C. P. (2000). Toward a deconstruction of the metaphor of behavioral momentum. *Behavioral and Brain Sciences*, 23, 111–112.
- Shull, R. L., Gaynor, S. T., & Grimes, J. A. (2001). Response rate viewed as engagement bouts: Effects of relative reinforcement and schedule type. *Journal of the Experimental Analysis of Behavior*, 75, 247–274.
- Sibley, R. M., Nott, H. M. R., & Fletcher, D. J. (1990). Splitting behaviour into bouts. *Animal Behaviour*, 39, 63–69.
- Tolkamp, B. J., Allcroft, D. J., Austin, E. J., Nielsen, B. L., & Kyriazakis, I. (1998). Satiety splits feeding behaviour into bouts. *Journal of Theoretical Biology*, 194, 235–250.
- Walter, D. E., & Palya, W. L. (1984). An inexpensive experiment controller for stand-alone applications or distributed processing networks. *Behavior Research Methods, Instruments, & Computers*, 16, 125–134.

Received July 28, 2001

Final acceptance January 5, 2002

APPENDIX

The performance measures from which the ratios (logarithms) shown in Figure 3 were derived. The third through the eighth columns from the left list the three-session means from the baseline sessions immediately before the extinction session. (For Condition 2b, Rats A1 through B3, the entries indicate performance on just the last session before the extinction session because IRT data were not obtained from the earlier two sessions.) These data were the basis of the ratios in the top three panels of Figure 3. The four columns on the right show the data from which the relative-resistance-to-extinction measure was computed (bottom panel of Figure 3). Baseline response rates were means taken over the last five sessions before the extinction session. Condition 1 arranged a VI schedule in both components of the multiple schedule (VI 4 min vs. VI 1 min). Condition 2 arranged equal VI schedules in the two components of the multiple schedule but added a response-independent (VT) schedule of pellet delivery in one of the components (VI 4 min vs. VI 4 min + VT 1.33 min). Condition 2b is a replication of Condition 2a, following additional training.

Con- di- tion	Rat	Mean of last three sessions before extinction						Data for assessing resistance to extinction			
		Total responses per minute		Bouts per minute		Responses per bout		Baseline five-session mean responses per minute		Extinction session responses per minute	
		Rich	Lean	Rich	Lean	Rich	Lean	Rich	Lean	Rich	Lean
1	11	47.36	24.38	14.24	8.19	4.69	3.89	43.60	22.06	12.80	4.32
	12	60.24	35.17	14.16	11.52	7.25	4.07	69.40	43.69	10.60	1.44
	13	40.63	21.35	8.99	5.41	6.50	5.46	55.21	24.23	13.24	1.68
	14	16.46	6.23	11.32	5.34	2.01	1.56	16.12	7.03	3.12	0.04
	A1	26.64	11.88	10.34	5.92	3.63	3.06	29.97	11.11	6.20	0.76
	A2	17.87	8.89	10.03	4.74	2.00	2.08	17.81	10.06	7.88	3.68
	A3	10.97	5.64	6.19	3.07	3.08	3.08	10.30	5.75	6.64	1.44
	A4	46.57	19.63	9.64	7.95	7.00	5.75	49.47	19.10	14.68	4.08
	B1	49.45	31.19	13.45	10.33	4.83	3.92	48.27	30.03	31.28	11.88
	B2	36.03	27.04	10.14	8.77	4.81	3.91	32.41	22.96	21.28	8.32
2a	B3	16.27	9.35	9.37	6.73	2.39	2.23	16.74	8.95	9.28	3.32
	11	29.68	26.64	7.11	7.10	5.18	4.97	29.86	24.35	11.36	5.40
	12	21.49	22.12	14.16	11.16	2.19	2.44	22.37	22.21	9.16	2.80
	13	11.93	9.87	4.15	2.33	3.42	3.78	12.10	9.70	9.16	3.32
	14	3.73	3.71	4.66	3.79	1.04	1.14	3.94	3.61	2.28	0.48
	A1	16.98	9.75	8.88	5.14	2.61	2.82	16.04	10.52	13.36	7.24
	A2	11.05	11.90	6.08	6.94	2.19	2.47	11.33	11.55	4.44	3.40
	A3	6.34	5.66	4.01	4.51	2.90	2.95	6.05	5.06	2.24	0.92
	A4	20.32	14.82	6.52	7.20	4.12	3.91	21.18	14.93	8.76	5.96
	B1	31.43	26.44	7.00	6.63	4.82	4.21	29.66	24.33	14.36	8.68
2b	B2	11.80	10.56	5.96	3.63	3.00	3.00	14.00	12.13	8.52	3.24
	B3	9.06	6.30	3.87	4.07	2.95	1.99	8.80	5.59	5.88	3.40
	11	18.30	15.00	3.85	4.95	5.83	5.08	16.48	13.82	5.48	2.36
	12	18.73	20.19	8.66	10.73	3.18	2.39	18.71	18.96	12.12	6.08
	13	12.65	12.84	4.59	3.45	3.71	4.10	12.25	13.17	7.80	2.52
	14	2.99	2.65	3.67	1.97	1.04	1.11	3.14	2.92	1.68	0.84
	A1	19.74	12.60	13.25	6.58	2.06	2.39	18.91	14.67	8.20	3.52
	A2	11.70	11.55	4.46	6.57	2.75	2.46	11.10	11.80	4.32	3.80
	A3	4.50	2.67	2.71	2.47	2.49	2.38	3.89	2.96	2.40	0.40
	A4	12.96	13.62	3.88	7.82	3.36	3.51	19.42	16.24	9.04	4.12
	B1	26.78	23.34	6.20	77.83	4.50	3.15	26.61	23.02	12.68	7.68
	B2	12.78	13.23	5.02	3.41	3.45	3.75	13.12	13.22	9.24	3.64
	B3	7.86	8.10	1.85	4.27	3.58	2.48	8.16	8.46	6.16	3.24